


IN THE UNITED STATES PATENT AND TRADEMARK OFFICE BEFORE THE
BOARD OF PATENT APPEALS AND INTERFERENCES

Appellant: Simon Knowles
Serial No.: 10/813,615
Filed: March 31, 2004
Title: APPARATUS AND METHOD FOR ASYMMETRIC DUAL PATH
PROCESSING
Grp./A.U.: 2183
Examiner: Robert E. Fennema
Confirmation No.: 3818

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Mail Stop Appeal Brief-Patents



ATTENTION: Board of Patent Appeals and Interferences

Sirs:

APPEAL BRIEF UNDER 37 C.F.R. § 41.37

This is an appeal from a Final Rejection dated August 4, 2009, of Claims 1-6, 8-18 and 21. The Appellant submits this Brief with the statutory fee of \$540.00 as set forth in 37 C.F.R. § 41.20(b)(2), and hereby authorize the Commissioner to charge any additional fees connected with this communication or credit any overpayment to Deposit Account No. 08-2395.

This Brief contains these items under the following headings, and in the order set forth below in accordance with 37 C.F.R. §41.37(c)(1):

- I. REAL PARTY IN INTEREST
- II. RELATED APPEALS AND INTERFERENCES
- III. STATUS OF CLAIMS
- IV. STATUS OF AMENDMENTS
- V. SUMMARY OF CLAIMED SUBJECT MATTER
- VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL
- VII. APPELLANTS' ARGUMENTS
- VIII. APPENDIX A – CLAIMS
- IX. APPENDIX B – EVIDENCE
- X. RELATED PROCEEDINGS APPENDIX

I. REAL PARTY IN INTEREST

The real party in interest in this appeal is the assignee, Icera Inc.

II. RELATED APPEALS AND INTERFERENCES

Appellant does not know of any prior or pending Appeals, Interferences, or Judicial Proceedings directly related to, affecting, affected by, or have a bearing on the Board's decision in this appeal.

III. STATUS OF THE CLAIMS

Claims 1-6, 8-18 and 21 stand rejected.

Herein all rejections of Claims 1-6, 8-18 and 21 are being appealed.

IV. STATUS OF THE AMENDMENTS

An amendment to the Final Rejection was filed on October 5, 2009. In the amendment, Claims 1, 8, 11, 18 and 21 were amended and Claim 7 was canceled without prejudice or disclaimer. Claims 19-20 were canceled in a previous amendment.

The Advisory Action dated October 22, 2009, indicated that the amendment of October 5, 2009, would be entered for the purpose of an appeal. These Claims are presented in Section VIII, Appendix A of this Appeal Brief. No other amendments are pending.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Independent Claim 1 is directed to a computer processor for processing: (i) instruction packets including a plurality of only control instructions, the control instructions having a control bit width, and (ii) instruction packets including a plurality of instructions including at least one data processing instruction, the data processing instructions having a data processing bit width wider than the control bit width. The processor includes: (1) a decode unit for decoding sequentially the instruction packets fetched from a memory holding the instruction packets, the instruction packets being all of equal bit length, (2) a control processing channel capable of performing control operations, the control processing channel including a plurality of functional units including a control register file having a first bit width and (3) a data processing channel capable of performing data processing operations at least one input of which is a vector, the data processing channel including a plurality of functional units including a data register file having a second bit width, wider than the first bit width. The decode unit includes decode circuitry configured to decode identification bits of each instruction packet to determine which type (i), (ii), of instruction packet is being decoded, and control circuitry configured to pass the plurality of only control instructions having the control bit width from an instruction packet of type (i) to the control processing channel when the decode circuitry indicates so and to pass the plurality of instructions including at least one data processing instruction having the data processing bit width wider than the control bit width from an instruction packet of type (ii) to the data processing channel when the decode circuitry indicates so. In use the decode unit causes instructions of (i) instruction packets including a plurality of only control instructions to be executed sequentially on the control processing channel and, in use the decode unit causes instructions of (ii) instruction packets including a plurality of instructions including at least one data

processing instruction to be executed simultaneously on the data processing channel. (See page 2, lines 2-13.)

Independent Claim 18 is directed to method of operating a computer processor for processing: (i) instruction packets including a plurality of only control instructions, the control instructions having a control bit width, and (ii) instruction packets including a plurality of instructions including at least one data processing instruction, the data processing instructions having a data processing bit width wider than the control bit width. The processor includes: (1) a decode unit for decoding sequentially the instruction packets fetched from a memory holding the instruction packets, the instruction packets being all of equal bit length; (2) a control processing channel including a plurality of functional units including a control register file having a first bit width; and (3) a data processing channel capable of performing data processing operations at least one input of which is a vector, the data processing channel including a plurality of functional units including a data register file having a second bit width, wider than the first bit width. The method including: (1) decoding identification bits of each instruction packet to determine which type (i), (ii), of instruction packet is being decoded, and passing the plurality of only control instructions having the control bit width from an instruction packet of type (i) to the control processing channel when the decode circuitry indicates so and passing the plurality of instructions including at least one data processing instruction having the data processing bit width wider than the control bit width from an instruction packet of type (ii) to the data processing channel when the decode circuitry indicates so; (2) supplying, when the instruction packet defines (i) a plurality of only control instructions, the control instructions to the control processing channel wherein the control instructions are executed sequentially; and (3) supplying, when the instruction packet defines (ii) a plurality of instructions including at least one data processing instruction, at least the data instruction to the data processing

channel wherein the plurality of instructions are executed simultaneously. (See page 3, line 18, to page 4, line 7.)

Independent Claim 21 is directed to a computer readable medium including a sequence of instruction packets, the instruction packets being all of equal bit length. The instruction packets include a first type of instruction packet including a plurality of only control instructions of equal width, the control instructions having a control bit width, and a second type of instruction packet including a plurality of instructions including at least one data processing instruction, the at least one data processing instructions having a data processing bit width wider than the control bit width, and wherein at least one data processing instruction is a vector. The instruction packets include at least one indicator bit at a designated bit location within the instruction packet, wherein the computer readable-medium is adapted to run on a computer such that said indication bit is adapted to cooperate with a decode unit of the computer to designate whether: (1) the instruction packet defines a plurality of only control instructions having the control bit width or a plurality of instructions including at least one data processing instruction having the data processing bit width wider than the control bit width and (2) in the case when there is a plurality of instructions including at least one data instruction, the nature of each of the first and second instructions selected from: a control instruction; a data instruction; and a memory access instruction. (See page 5, lines 8-19.)

In one embodiment, for example, the original specification discloses an asymmetric dual path computer processor in FIG. 1. The processor of FIG. 1 divides processing of a single instruction stream 100 between two different hardware execution paths: a control execution path 102, which is dedicated to processing control code, and a data execution path 103, which is dedicated to processing data code. The data widths, operators, and other characteristics of the two execution paths 102, 103

differ according to the different characteristics of control code and datapath code. In the processor of FIG. 1, the two different execution paths 102 and 103 are dedicated to handling the two different types of code, with each side having its own architectural register file, such as control register file 104 and data register file 105, differentiated by width and number of registers; the control registers are of narrower width, by number of bits, and the data registers are of wider width. The processor is therefore asymmetric, in that its two execution paths are different bit-widths owing to the fact that they each perform different, specialized functions. (See page 6, line 14, to page 7, line 7.)

In addition to the control execution path 102 and the data execution path 103, the processor of FIG. 1 includes an instruction decode unit 101. (See page 7, line 8, to page 8, line 9.) In one embodiment, the instruction decode unit 101 decodes identification bits of each instruction packet to determine which type of packet is being decoded. Having decoded the initial bits of each instruction packet, the instructions of each packet are passed to either the control execution path 102 or the data execution path 103 according to the type of instruction. (See page 10, lines 11-19; page 6, line 14, to page 7, line 7; page 7, line 8, to page 8, line 9; and FIG. 1.)

Thus, in the disclosed embodiment of the FIG. 1, the processor divides processing of a single instruction stream 100 into two different hardware execution paths; a control execution path 102 and a data execution path 103. (See page 6, line 15, to page 7, line 7.) The instruction stream 100 is made up of a series of instruction packets. As an example, FIG. 2 of the original specification shows three types of instruction packets for the processor of FIG. 1. Instruction packet 211 is a 3-scalar type, for dense control code, and includes three 21-bit control instructions (c21). Instruction packets 212 and 213 are LIW (long instruction word) type, for parallel execution of datapath code. In this example each instruction packet 212, 213 includes two instructions but different numbers may be included if desired. (See page 9, line 21, to page 10, line 3.)

Instruction decode unit 101 of the embodiment of FIG. 1 uses the initial identification bits, or some other designated identification bits at predetermined bit locations, of each instruction packet to determine which type of packet is being decoded. For example, as shown in FIG. 2, an initial indicator bit "1" signifies that an instruction packet is of a scalar control instruction type, with three control instructions; while initial indicator bits "0 1" and "0 0" signify instruction packets of type 212 and 213, with a data and memory instruction in packet 212 or a data and control instruction in packet 213. (See page 10, lines 11-17.) In order to execute the instruction packets represented in FIG. 2, the instruction decode unit 101 fetches program packets from memory. (See page 10, lines 20-21.)

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

(A) Whether Claim 1 is obvious over the combination of U.S. Patent No. 5,922,065 to Hull ("Hull"), in view of an article entitled, "Unifying FPGAs and SIMD Arrays" by Bolotski, *et al.* ("Bolotski") and further in view of an article entitled, "Computer Architecture: A Quantitative Approach" by Hennessy ("Hennessy") as applied by the Final Rejection at pages 3-6.

(B) Whether Claim 18 is obvious over the combination of Hull, Bolotski and Hennessy as applied by the Office Action at pages 9-12.

(C) Whether Claim 21 is obvious over the combination of Hull, Bolotski and Hennessy as applied by the Office Action at pages 12-14.

(D) Whether Claims 2-6 and 14-17 are obvious over the combination of Hull, Bolotski and Hennessy as applied by the Office Action at pages 6-9.

(E) Whether Claims 8-10 and 12-13 are obvious over the combination of Hull, Bolotski, Hennessy and *In Re Rose* as applied by the Office Action at pages 14-15.

VII. APPELLANT'S ARGUMENT

The inventions set forth in independent Claims 1, 18 and 21, and their respective dependent claims are not obvious over the applied references on which the Examiner relies.

(A) Regarding the Grounds of Rejection (A), the obviousness rejection of Claim 1 is improper.

(1) The obviousness rejection is improper because it relies on Hull to teach features that are not taught in the sections of Hull that are relied upon.

In general, Hull fails to teach or suggest a processor having a decode unit, a control processing channel and a data processing channel as recited in Claim 1, wherein the decode unit determines to pass an instruction packet to either the control processing channel or the data processing channel for processing after determining the type of packet. (See, for example, FIG. 1 of the present specification as an example of an embodiment of a disclosed processor.) Instead of decoding, Hull relates to a processor used for instruction encoding. (See column 1, lines 7-9.) Unlike Claim 1, Hull relates to encoding instruction sequences that identify each instruction to correspond to execution units of the processor. (See column 2, lines 11-21.) A template field is used to map instruction slots to corresponding execution units. (See column 2, lines 22-27.) Thus, unlike Claim 1, Hull relates to encoding instructions to improve efficiencies of existing processor architectures by directing types of instructions to corresponding types of execution units. (See, for example, column 2, lines 22-29.)

More specifically, Hull as applied fails to teach or suggest each limitation of Claim 1 as relied upon in the rejection. This includes the decode unit, the control processing channel and the data processing channel of Claim 1. For example, the Examiner relies on column 5, lines 16-18, of Hull to teach a decode unit. (See Final Rejection, page 3.) This cited portion of Hull does not

disclose or suggest a decode unit but states: “(I)nstructions in bundles with lower memory addresses are considered to precede instructions in bundles with higher memory addresses.”

Similarly, Hull as applied also fails to teach or suggest the control processing channel and the data processing channel. The cited portions relied on of Hull merely disclose different execution units; not processing channels as included in the processor architecture of Claim 1. (See Final Rejection, pages 3-4.) Accordingly, Hull fails to teach or suggest teach limitation for which it has been relied upon.

Additionally, the Examiner appears to rely on the template field of Hull to disclose decoding identification bits of each instruction packet to determine which type of instruction packet is being decoded. (See Final Rejection, page 4, referring to column 3, lines 63-66, of Hull.) The template field of Hull, however, is used for mapping instruction slots to execution type units. (See column 3, lines 65-66.) An instruction slot is not an instruction packet as recited in Claim 1 but instead is part of an instruction bundle that includes multiple instruction slots. (See column 3, lines 52-60 and FIG. 3.) Thus, instead of determining which type of instruction packet is being decoded as recited in Claim 1, Hull is concerned with the individual instructions in a bundle and mapping the individual instructions to specific execution units. Hull, therefore, operates with a different architecture that does not appear to be concerned with which channel to send an instruction packet for processing but with what execution unit to send an instruction from an instruction slot. As such, Hull also fails to teach or suggest this limitation for which it has been relied upon.

Furthermore, the Examiner relies on column 4, lines 61-62 of Hull to disclose “wherein, in use the decode unit causes instructions of (i) instruction packets comprising a plurality of only control instructions to be executed sequentially on the control processing channel.” (See Final Rejection, page 4.) Column 4, lines 61-65 of Hull state: “Within a bundle, execution

order proceeds from slot 0 to slot 2. If the S-bit is 0, the instruction group containing the last instruction (slot 2) of the current bundle continues into the first instruction (slot 0) of the statically next sequential bundle.” As noted in column 3, lines 24-27 of Hull, the term “bundle” refers to three instructions and a template field. The cited portion of Hull, therefore, discloses execution orders for individual instructions of an instruction bundle to proceed from slot 0 to slot 2, but fails to teach or suggest where “instruction packets” that have a plurality of only control instructions are caused to be executed sequentially on the control processing channel. As such, Hull fails to teach or suggest this feature as relied upon.

The Examiner also relies on column 2, lines 5-9 of Hull to teach: “wherein, in use the decode unit causes instructions of (ii) instruction packets comprising a plurality of instructions comprising at least one data processing instruction to be executed simultaneously on the data processing channel.” (See Final Rejection, page 4.) Column 2, lines 5-9 of Hull states: “As will be seen, the present invention provides a processor capable of simultaneously executing a plurality of sequential instructions with a highly-efficient encoding of instructions.” This cited portion provides no teaching or even suggestion: (1) a decode unit, (2) causing instructions including at least one data processing instruction, (3) to be executed simultaneously on a data processing channel. As such, Hull fails to teach or suggest this feature as relied upon.

Neither Bolotski nor Hennessy have been cited to cure the above noted deficiency of Hull but to address other deficiencies of Hull noticed by the Examiner. Due to the above-described deficiencies, the Office Action or Advisory Action does not cite prior art teachings for each feature in pending independent Claim 1. For that reason, the rejection of pending independent Claim 1 does not provide a *prima facie* case of obviousness. Accordingly, the Appellant respectfully requests the Appeal Board to reverse the rejection of the Examiner and allow issuance of Claim 1.

(2) The obviousness rejection is improper because it makes an improper combination of Hennessy and Hull.

The Examiner asserts that Hull "does not teach a variable-length instruction set, and thus, cannot teach a control instruction with a bit-width shorter than a data processing instruction." (*See* Final Rejection, page 15, point 21.) The Examiner applies Hennessy to teach there are different ways to teach an instruction set (variable, fixed or a hybrid of the two) and each way has advantages or disadvantages. (*See* Final Rejection, page 16, point 21.)

One skilled in the art, however, would not be motivated to apply the asserted teachings of Hull to a variable-length instruction set as asserted by the Examiner. (*See*, for example, the Advisory Action.) On the contrary, Hull is directed to reducing waste and inefficiency in encoding associated with fixed formats. (*See* column 2, lines 3-8, and column 5, lines 5-8.) Hull discloses using a template field to specify group boundaries within a bundle and the mapping of instruction slots to execution unit types. (*See*, for example, column 4, lines 20-26 and FIG. 4.) One skilled in the art, however, would not be motivated to employ the teachings of Hull in a variable-length format since a variable-length format would not have the waste and inefficiency of encoding that Hull addresses. In other words, Hull addresses problems associated with a fixed-length format. As such, one skilled in the art would not find Hull beneficial for a variable-length architecture. Accordingly, the combination of Hennessy with Hull is improper.

Additionally, since Hull is for fixed-length formats and the combination of Hennessy with Hull is improper, the applied combination also fails to teach or suggest "control instructions having a control bit width" and "data processing instructions having a data processing bit width wider than the control bit width" as recited in Claim 1.

It is submitted, therefore, that there is nothing to suggest that a skilled person would have any motivation to consider Hennessy with Hull. In fact, the need for the Final Rejection to combine Hull with the textbook Hennessy to assert prior art teachings for some recited features of independent Claim 1 further indicates that Claim 1 is non-obvious. Thus, the rejection does not provide a *prima facie* case of obviousness of Claim 1. Accordingly, the Appellant respectfully requests the Appeal Board to reverse the rejection of Claim 1 and allow issuance thereof.

(B) Claim 18 is not obvious over the applied combination of Hull, Bolotski and Hennessy.

(1) Hull fails to teach or suggest each limitation upon which it has been relied.

In general, Hull fails to teach or suggest a method of operating a computer processor that includes a decode unit, a control processing channel and a data processing channel. Instead of decoding, Hull relates to a processor used for instruction encoding. (See column 1, lines 7-9.) Unlike Claim 18, Hull relates to encoding instruction sequences that identify each instruction to correspond to execution units of the processor. (See column 2, lines 11-21.) A template field is used to map instruction slots to corresponding execution units. (See column 2, lines 22-27.) Thus, unlike Claim 18, Hull relates to encoding instructions to improve efficiencies of existing processor architectures by directing types of instructions to corresponding types of execution units. (See, for example, column 2, lines 22-29.)

More specifically, Hull as applied fails to teach or suggest a method of operating a processor as recited in Claim 18 wherein the processor includes a decode unit, a control processing channel and a data processing channel. For example, the Examiner relies on column 5, lines 16-18, of Hull to teach a decode unit. (See Final Rejection, page 10.) This cited portion of Hull does not disclose or

suggest a decode unit but states: “(I)nstructions in bundles with lower memory addresses are considered to precede instructions in bundles with higher memory addresses.”

Similarly, Hull as applied also fails to teach or suggest the control processing channel and the data processing channel. The cited portions relied on of Hull merely disclose different execution units; not processing channels as included in the processor architecture of Claim 18. (See Final Rejection, page 10.) Accordingly, Hull fails to teach or suggest teach limitation for which it has been relied upon.

Additionally, the Examiner appears to rely on the template field of Hull to disclose decoding identification bits of each instruction packet to determine which type of instruction packet is being decoded. (See Final Rejection, pages 10-11, referring to column 3, lines 63-66, of Hull.) The template field of Hull, however, is used for mapping instruction slots to execution type units. (See column 3, lines 65-66.) An instruction slot is not an instruction packet as presently claimed but instead is part of an instruction bundle that includes multiple instruction slots. (See column 3, lines 52-60 and FIG. 3.) Thus, instead of determining which type of instruction packet is being decoded as recited in Claim 18, Hull is concerned with the individual instructions in a bundle and mapping the individual instructions to specific execution units. Hull, therefore, operates with a different architecture that does not appear to be concerned with which channel to send an instruction packet for processing but with what execution unit to send an instruction from an instruction slot. As such, Hull also fails to teach or suggest this limitation for which it has been relied upon.

Furthermore, the Examiner relies on column 4, lines 61-62 of Hull to disclose “when the instruction packet defines (i) a plurality of only control instructions ... wherein the control instructions are executed sequentially.” (See Final Rejection, page 11.) Column 4, lines 61-65 of Hull state: “Within a bundle, execution order proceeds from slot 0 to slot 2. If the S-bit

is 0, the instruction group containing the last instruction (slot 2) of the current bundle continues into the first instruction (slot 0) of the statically next sequential bundle.” As noted in column 3, lines 24-27 of Hull, the term “bundle” refers to three instructions and a template field. The cited portion of Hull, therefore, discloses execution orders for individual instructions of an instruction bundle to proceed from slot 0 to slot 2, but fails to teach or suggest where “instruction packets” that have a plurality of only control instructions are executed sequentially on a control processing channel. As such, Hull fails to teach or suggest this feature as relied upon.

The Examiner also relies on column 2, lines 5-9 of Hull to teach: “when the instruction packet defines (ii) a plurality of instructions comprising at least one data processing instruction ... to the data processing channel wherein the plurality of instructions are executed simultaneously.” (See Final Rejection, page 11.) Column 2, lines 5-9 of Hull states: “As will be seen, the present invention provides a processor capable of simultaneously executing a plurality of sequential instructions with a highly-efficient encoding of instructions.” This cited portion provides no teaching or even suggestion of: (1) causing instructions including at least one data processing instruction, (2) to be executed simultaneously on a data processing channel. As such, Hull fails to teach or suggest this feature as relied upon.

Neither Bolotski nor Hennessy have been cited to cure the above noted deficiency of Hull but to address other deficiencies of Hull noticed by the Examiner. Due to the above-described deficiencies, the Office Action or Advisory Action does not cite prior art teachings for each feature in pending independent Claim 18. For that reason, the rejection of pending independent Claim 1 does not provide a *prima facie* case of obviousness. Accordingly, the Appellant respectfully requests the Appeal Board to reverse the rejection of the Examiner and allow issuance of Claim 18.

(2) The combination of Hennessy with Hull is improper.

The Examiner asserts that Hull "does not teach a variable-length instruction set, and thus, cannot teach a control instruction with a bit-width shorter than a data processing instruction." (*See* Final Rejection, page 15, point 21.) The Examiner applies Hennessy to teach there are different ways to teach an instruction set (variable, fixed or a hybrid of the two) and each way has advantages or disadvantages. (*See* Final Rejection, page 16, point 21.)

One skilled in the art, however, would not be motivated to apply the asserted teachings of Hull to a variable-length instruction set as asserted by the Examiner. (*See*, for example, the Advisory Action.) On the contrary, Hull is directed to reducing waste and inefficiency in encoding associated with fixed formats. (*See* column 2, lines 3-8, and column 5, lines 5-8.) Hull discloses using a template field to specify group boundaries within a bundle and the mapping of instruction slots to execution unit types. (*See*, for example, column 4, lines 20-26 and FIG. 4.) One skilled in the art, however, would not be motivated to employ the teachings of Hull in a variable-length format since a variable-length format would not have the waste and inefficiency of encoding that Hull addresses. In other words, Hull addresses problems associated with a fixed-length format. As such, one skilled in the art would not find Hull beneficial for a variable-length architecture. Accordingly, the combination of Hennessy with Hull is improper.

Additionally, since Hull is for fixed-length formats and the combination of Hennessy with Hull is improper, the applied combination also fails to teach or suggest "control instructions having a control bit width" and "data processing instructions having a data processing bit width wider than the control bit width" as recited in Claim 18.

Thus, for at least the above reasons, the applied combination of Hull, Bolotski and Hennessy does not provide a *prima facie* case of obviousness of independent Claim 18. Accordingly, the

Appellant respectfully requests the Appeal Board to reverse the rejection of Claim 18 and allow issuance thereof.

(C) Claim 21 is not obvious over the applied combination of Hull, Bolotski and Hennessy.

(1) The combination of Hennessy with Hull is improper.

The Examiner asserts Hull and Bolostki are silent towards data processing instructions being wider than a control instruction, because Hull teaches a fixed-length architecture. (*See* Final Rejection, page 14.) The Examiner also asserts that Hull "does not teach a variable-length instruction set, and thus, cannot teach a control instruction with a bit-width shorter than a data processing instruction." (*See* Final Rejection, page 15, point 21.) The Examiner applies Hennessy to teach there are different ways to teach an instruction set (variable, fixed or a hybrid of the two) and each way has advantages or disadvantages. (*See* Final Rejection, page 16, point 21.)

One skilled in the art, however, would not be motivated to apply the asserted teachings of Hull to a variable-length instruction set as asserted by the Examiner. (*See*, for example, the Advisory Action.) On the contrary, Hull is directed to reducing waste and inefficiency in encoding associated with fixed formats. (*See* column 2, lines 3-8, and column 5, lines 5-8.) Hull discloses using a template field to specify group boundaries within a bundle and the mapping of instruction slots to execution unit types. (*See*, for example, column 4, lines 20-26 and FIG. 4.) One skilled in the art, however, would not be motivated to employ the teachings of Hull in a variable-length format since a variable-length format would not have the waste and inefficiency of encoding that Hull addresses. In other words, Hull addresses problems associated with a fixed-length format. As such, one skilled in the art would not find Hull beneficial for a variable-length architecture. Accordingly, the combination of Hennessy with Hull is improper.

Thus, for at least the above reasons, including the above arguments addressing the rejections of Claims 1 and 18, the applied combination of Hull, Bolotski and Hennessy does not provide a *prima facie* case of obviousness of independent Claim 21. Accordingly, the Appellants respectfully request the Appeal Board to reverse the rejection of Claim 21 and allow issuance thereof.

(D) Claims 2-6 and 14-17 are not obvious over the applied combination of Hull, Bolotski and Hennessy.

Claims 2-7 and 14-17 are non-obvious over the above combination, as applied by the Final Rejection, at least by their dependence on independent Claim 1.

(E) Claims 8-10 and 12-13 are not obvious over the applied combination of Hull, Bolotski, Hennessy and *In Re Rose*.

Claims 8-10 and 12-13 are non-obvious over the above combination, as applied by the Final Rejection, at least by their dependence on independent Claim 1.

CONCLUSION

For the reasons set forth above, the Claims on appeal are patentably non-obvious over the applied references. Accordingly, the Appellant respectfully requests that the Board of Patent Appeals and Interferences reverse the Final Rejection of all of the pending claims and allow issuance thereof.

Respectfully submitted,

HITT GAINES, P.C.

/J. Joel Justiss/

J. Joel Justiss
Registration No. 48,981

Dated: March 1, 2010

Hitt Gaines, PC
P. O. Box 832570
Richardson, Texas 75083-2570
(972) 480-8800
(972) 480-8865 (Fax)
joel.justiss@hittgaines.com

VIII. APPENDIX A – CLAIMS

1. (Previously Presented) A computer processor for processing (i) instruction packets comprising a plurality of only control instructions, the control instructions having a control bit width, and (ii) instruction packets comprising a plurality of instructions comprising at least one data processing instruction, the data processing instructions having a data processing bit width wider than the control bit width, the processor comprising:

a decode unit for decoding sequentially the instruction packets fetched from a memory holding the instruction packets, the instruction packets being all of equal bit length;

a control processing channel capable of performing control operations, the control processing channel comprising a plurality of functional units including a control register file having a first bit width; and

a data processing channel capable of performing data processing operations at least one input of which is a vector, the data processing channel comprising a plurality of functional units including a data register file having a second bit width, wider than the first bit width;

wherein the decode unit comprises decode circuitry configured to decode identification bits of each instruction packet to determine which type (i), (ii), of instruction packet is being decoded, and control circuitry configured to pass the plurality of only control instructions having the control bit width from an instruction packet of type (i) to the control processing channel when the decode circuitry indicates so and to pass the plurality of instructions comprising at least one data processing instruction having the data processing bit width wider than the control bit width from an instruction packet of type (ii) to the data processing channel when the decode circuitry indicates so;

wherein, in use the decode unit causes instructions of (i) instruction packets comprising a plurality of only control instructions to be executed sequentially on the control processing channel; and

wherein, in use the decode unit causes instructions of (ii) instruction packets comprising a plurality of instructions comprising at least one data processing instruction to be executed simultaneously on the data processing channel.

2. (Previously Presented) A computer processor according to claim 1, wherein the control processing channel further comprises a branch unit and a control execution unit.

3. (Previously Presented) A computer processor according to claim 1, wherein the data processing channel further comprises a fixed data execution unit and a configurable data execution unit.

4. (Original) A computer processor according to claim 3, wherein the fixed data execution unit and the configurable data execution unit both operate according to a single instruction multiple data format.

5. (Previously Presented) A computer processor according to claim 1, wherein the control and data processing channels share a load store unit.

6. (Previously Presented) A computer processor according to claim 5, wherein the load store unit uses control information supplied by the control processing channel and data supplied by the data processing channel.

7. (Canceled)

8. (Previously Presented) A computer processor according to claim 1, wherein the instruction packets are all of a 64-bit length.

9. (Original) A computer processor according to claim 1, wherein the control instructions are all of a bit length between 18 and 24 bits.

10. (Original) A computer processor according to claim 9, wherein the control instructions are all of a 21-bit length.

11. (Previously Presented) A computer processor according to claim 1, wherein the nature of each instruction in an instruction packet is selected at least from a control instruction, a data instruction, and a memory access instruction.

12. (Original) A computer processor according to claim 11, wherein the bit length of each data instruction is 34 bits.

13. (Original) A computer processor according to claim 11, wherein the bit length of each memory access instruction is 28 bits.

14. (Previously Presented) A computer processor according to claim 1, wherein when the decode unit detects that the instruction packet defines three control instructions, the decode unit is operable to supply the control processing channel with the three control instructions whereby the three control instructions are executed sequentially.

15. (Previously Presented) A computer processor according to claim 1, wherein when the decode unit detects that the instruction packet defines two instructions comprising at least one data instruction, the decode unit is operable to supply the data processing channel with at least the data

instruction whereby the two instructions are executed simultaneously.

16. (Previously Presented) A computer processor according to claim 1, wherein the decode unit is operable to read the values of a set of designated bits at predetermined bit locations in each instruction packet of the sequence, to determine:

a) whether the instruction packet defines a plurality of control instructions or a plurality of instructions of which at least one is a data instruction; and

b) where the instruction packet defines a plurality of instructions of which at least one is a data instruction, the nature of each of the two instructions selected from: a control instruction; a data instruction; and a memory access instruction.

17. (Original) A computer processor according to claim 3, wherein the configurable data execution unit is capable of executing more than two consecutive operations on the data provided by a single issued instruction before returning a result to a destination register file.

18. (Previously Presented) A method of operating a computer processor for processing (i) instruction packets comprising a plurality of only control instructions, the control instructions having a control bit width, and (ii) instruction packets comprising a plurality of instructions comprising at least one data processing instruction, the data processing instructions having a data processing bit width wider than the control bit width, the processor comprising a decode unit for decoding sequentially the instruction packets fetched from a memory holding the instruction packets, the instruction packets being all of equal bit length; a control processing channels comprising a plurality of functional units including a control register file having a first bit width; and a data processing channel capable of performing data

processing operations at least one input of which is a vector, the data processing channel comprising a plurality of functional units including a data register file having a second bit width, wider than the first bit width, the method comprising:

decoding identification bits of each instruction packet to determine which type (i), (ii), of instruction packet is being decoded, and passing the plurality of only control instructions having the control bit width from an instruction packet of type (i) to the control processing channel when the decode circuitry indicates so and passing the plurality of instructions comprising at least one data processing instruction having the data processing bit width wider than the control bit width from an instruction packet of type (ii) to the data processing channel when the decode circuitry indicates so;

when the instruction packet defines (i) a plurality of only control instructions supplying the control instructions to the control processing channel wherein the control instructions are executed sequentially; and

when the instruction packet defines (ii) a plurality of instructions comprising at least one data processing instruction, supplying at least the data instruction to the data processing channel wherein the plurality of instructions are executed simultaneously.

Claims 19-20. (Canceled)

21. (Previously Presented) A computer readable-medium comprising a sequence of instruction packets, the instruction packets being all of equal bit length,

said instruction packets including a first type of instruction packet comprising a plurality of only control instructions of equal width, the control instructions having a control bit width, and a second type of instruction packet comprising a plurality of instructions

comprising at least one data processing instruction, the at least one data processing instructions having a data processing bit width wider than the control bit width, and wherein at least one data processing instruction is a vector,

said instruction packets including at least one indicator bit at a designated bit location within the instruction packet, wherein the computer readable-medium is adapted to run on a computer such that said indication bit is adapted to cooperate with a decode unit of the computer to designate whether:

a) the instruction packet defines a plurality of only control instructions having the control bit width or a plurality of instructions comprising at least one data processing instruction having the data processing bit width wider than the control bit width; and

b) in the case when there is a plurality of instructions comprising at least one data instruction, the nature of each of the first and second instructions selected from: a control instruction; a data instruction; and a memory access instruction.

IX. APPENDIX B – EVIDENCE

The evidence in this appendix includes the cited U.S. Patent to Hull and the cited article by Bolotski. Portions of the text book by Hennessy and the case of *In re Rose* were also cited by Examiner and relied upon in the Final Rejection.

X. RELATED PROCEEDINGS APPENDIX

NONE



Transit Note #95

Unifying FPGAs and SIMD Arrays

Michael Bolotski, Andre DeHon, and Thomas F. Knight, Jr.

Original Issue: September, 1993

Last Updated: Tue Mar 8 15:43:06 EST 1994

Acknowledgments: This research is supported in part by the Advanced Research Projects Agency under contracts N00014-91-J-1698 and DABT63-92-C-0039.

Abstract:

Field-Programmable Gate Arrays (FPGAs) and Single-Instruction Multiple-Data (SIMD) processing arrays share many architectural features. In both architectures, an array of simple, fine-grained logic elements is employed to provide high-speed, customizable, bit-wise computation. In this paper, we present a unified computational array model which encompasses both FPGAs and SIMD arrays. Within this framework, we examine the differences and similarities between these array structures and touch upon techniques and lessons which can be transferred between the architectures. The unified model also exposes promising prospects for hybrid array architectures. We introduce the Dynamically Programmable Gate Array (DPGA) which combines the best features from FPGAs and SIMD arrays into a single array architecture.

Introduction

FPGA-based custom computing engines and massively parallel SIMD arrays [BRV93][Bol93][Gea91] have been demonstrated to provide supercomputer-class performance on some tasks at a tiny fraction of supercomputer cost. Both of these architectures consist of arrays of small yet numerous processing elements. This similarity is the key to understanding their surprisingly high performance: most of the silicon in the FPGA and SIMD chips is actively operating on data bits. SIMD machines achieve high utilization by massive data parallelism. FPGA machines achieve high utilization by task-specific hardware configuration and pipelining. In both cases, several thousand bits are transformed per cycle, compared to the 64 bits of a typical microprocessor. In this paper we show that the similarities extend considerably further, and indeed that the two architectures can be viewed under a common computing model.



Field-Programmable Gate Arrays (FPGAs) are widely used today to implement general-purpose logic. FPGAs are built from a moderately fine-grained array of simple logic functions. An FPGA array is customized by selecting the logical function which each array element (AE) performs and the interconnection pattern between AEs. Using multiple stages of logic and primitive state elements, these arrays are programmed to implement both sequential and combinational general-purpose logic. FPGAs are in wide use today for system customization and glue logic, low-volume application-specific designs,






and IC prototyping. A variety of FPGAs are commercially available from a number of vendors (e.g. Xilinx [Inc91], Actel [Act90], Atmel [Fur93], Lattice [Qui92]).

Single-Instruction Multiple-Data (SIMD) arrays are employed to realize high throughput on many regular, computationally intensive data processing applications. An array of simple, fine-grained computational units make up most SIMD arrays. Typically, the computational units are wired together through local, nearest-neighbor communications. On each clock cycle, an instruction is broadcast to all AEs, and each AE performs the indicated computation on its local data element. SIMD arrays are commonly used for algorithms requiring regular, *data-parallel* computations where identical operations must be performed on a large set of data. Typical applications for SIMD arrays include low-level vision and image processing, discrete particle simulation, database searches, and genetic sequence matching. NASA's MPP [Pot85], Thinking Machines' CM-2 [Hil85], and MasPar's MP-1 [Bla90] are early examples of large-scale SIMD arrays. Increased silicon area along with advanced packaging trends allow production of very high-performance, highly-integrated, SIMD arrays at reasonable costs [BBLC93], [Lea91].

Viewed at an abstract level, these two array architectures are very similar. Both employ a moderately fine-grained array of logic elements. Each logic element performs a simple logic function on some state and some inputs from the array and either updates its own state to record the results of its computation or shares the results with other elements in the array. Despite the similarity, the design of FPGAs and SIMD arrays has evolved along quite divergent paths.


In this paper, we introduce a computational array model encompassing both FPGAs and SIMD arrays. By highlighting the tradeoffs made in the model to arrive at FPGA or SIMD structures, we can better appreciate optimizations made in the engineering of high-performance computational arrays. Further, we can transfer lessons learned between various array architectures. Additionally, the unified model highlights potentially novel hybrid array architectures. We examine one such hybrid, the *Dynamically Programmable Gate Array* (DPGA), and show that it can subsume the role of traditional FPGAs or SIMD arrays.



In Section  we introduce the uniform computational array model and show how FPGAs and SIMD arrays relate to this model. In Section  we compare the manner in which FPGA and SIMD arrays solve problems by considering the mapping of computations from one architecture to the other. In


Section  we introduce the DPGA hybrid architecture and describe its benefits and costs. In Sections  through , we examine these arrays at a somewhat abstract level, ignoring for the most part, technology-specific optimizations during implementation. In Section  we discuss the ways in which computational arrays are optimized based on computing style and available technology. In Section  we review the key themes explored in this paper.

Unified Computational Array Model



A computational array is composed from a regular lattice of AEs along with interconnection resources linking AEs together. Abstractly, each AE performs a simple computation on its inputs to produce one

or more output bits (See Figure ) . The inputs come from local state or via communication channels from other AEs. The outputs are either stored to local state or are communicated to other AEs. An *instruction* is used to specify the computation performed by each AE. Instructions are typically also used to specify communication and state manipulation, although in this paper we focus primarily on the instruction as specifying the operation of the logic unit. In this simplified model, the transform from

input values to outputs is modeled as a lookup table addressed by the input value (See Figure ) . We can model the *instruction* as either the programming of the lookup table or as additional inputs to the lookup table (See Figure ) .



Computational Block for an Array Element (AE)



In the most general setting, we might wish to specify a different instruction for each AE on each computational cycle. Unfortunately, if we design such an array, the resources required for instruction distribution dominate the array geometry and the required instruction bandwidth is unmanageably large. In practice, the manageable instruction bandwidth limits the operation rate of the array. With a

element array where each element implements distinct functions, Equation shows the relation between computational cycle time, and instruction bandwidth,

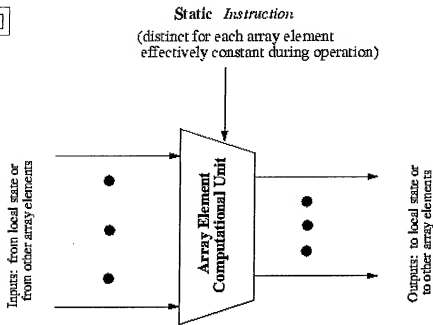
$$\frac{1}{N} \frac{1}{B}$$

For example, a 100 element array with 64-function elements operating at 10 MHz, requires $\frac{1}{100}$ ns = 6 Gbits/sec. If we are limited to an instruction distribution bandwidth of 1 Gbit/s, the clock cycle for the array must be limited to $\frac{1}{100} \frac{1}{1} = 600$ ns.

FPGAs and SIMD arrays both weaken this general computational model to avoid the requirement for huge instruction bandwidth. By simplifying the model, each type of array achieves a more pragmatic balance of resource requirements. For certain classes of applications these simplified models are adequate and can be engineered to take full advantage of the implementation technologies available.

FPGAs weaken the model by eliminating the instruction (See Figure) and therefore not changing each AE's operation through time. Different AEs, however, can be executing different operations. During a slow programming phase, each AE is configured with its operation, which remain fixed during subsequent normal operation. In SRAM-programmable FPGAs (e.g. Xilinx LCA [Inc91], Atmel [Eur93]) this programming phase normally occurs once each time the system is powered on. In fuse or anti-fuse based FPGAs (e.g. Actel [Act90]) a device is programmed exactly once during its lifetime.

$$\frac{1}{N}$$

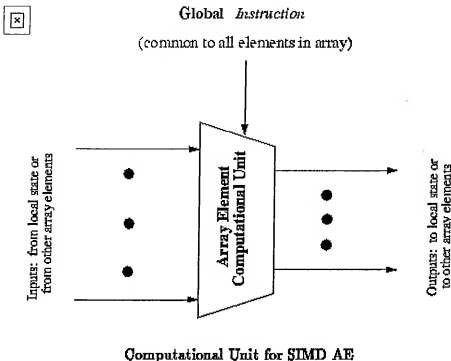


Computational Unit for FPGA AE



SIMD arrays weaken the model by distributing the *same* instruction to every AE (See Figure). Each AE is allowed to perform a different operation on each computational cycle, but all the elements in the array are required to perform the same operation during any given cycle.





These two compromises weaken the model along orthogonal dimensions. FPGAs compromise in the rate of instruction dispatch to allow the instructions to vary spatially through the array. SIMD arrays compromise in the spatial variation of instructions to allow a high rate of instruction dispatch.

Computational Equivalence x


We can view the compromises made in arriving at FPGA or SIMD arrays as space-time tradeoffs. FPGAs allow us to spatially construct any logical operation by composing an ensemble of AEs which implements the operation spatially. SIMD arrays allow us to temporally perform any logical operation by sequencing through the operations required to perform the operation in time. We can see this tradeoff further by considering how we can naively map an arbitrary FPGA computation onto a SIMD array and vice-versa.

SIMD Simulation on an FPGA

Since we can wire any sequential or combinational logic function on an FPGA array, we can simulate a SIMD AE by wiring up a group of FPGA AEs. Further, we can compose FPGA-implemented, SIMD AEs to simulate an arbitrary SIMD array. We can then run the same SIMD computation which we would have run on the SIMD array on the FPGA implementation.

Indeed, if the routing resources required to broadcast the same instruction to every simulated SIMD AE are directly available in the FPGA architecture, the implementation is extremely simple. In traditional SRAM-based FPGAs, one cell can serve as the SIMD AE computational element, one can serve as the local memory, and two can serve as the configurable interconnect.

FPGA Simulation on a SIMD Array

We can simulate any non-asynchronous FPGA in time using a SIMD array. To do so, we note that most SIMD arrays provide a bit mask to disable a subset of the cells for a particular operation. Conceptually, we map each FPGA AE to a corresponding SIMD AE. For each FPGA gate delay, we dispatch 


instructions, one for each different instruction the FPGA AE can perform. If necessary, we then dispatch a series of instructions to route output bits from producers to consumers. With each dispatch, the SIMD AEs are masked such that the only elements which perform a given operation are those which correspond to FPGA AEs which are programmed to perform the given operation.


Comparison


The FPGA implements the SIMD array spatially, and simulation overhead appears as a requirement for more gates and wiring resources. The SIMD array implements the FPGA temporally, and simulation overhead slows the computation rate. Of course, simulation is not the best approach for porting an algorithm from one architecture to the other. Nonetheless, the constructions above establish the feasibility of translating arbitrary computations between the two kinds of arrays.

Hybrid Architecture


In this section we introduce a hybrid architecture which allows instructions to vary both spatially and temporally without requiring additional instruction bandwidth.

As we saw in Figure , the lookup table modeling each SIMD, computational AE is programmed


identically on all AEs in the array. In Figure , we saw that each FPGA AE lookup table is programmed differently to select a different instruction on each AE. In a *Dynamically Programmable Gate Array* (DPGA) we allow the lookup table in each AE to be programmed differently. This allows

each AE to perform a different *operation* in response to each broadcast *instruction* (See Figure )

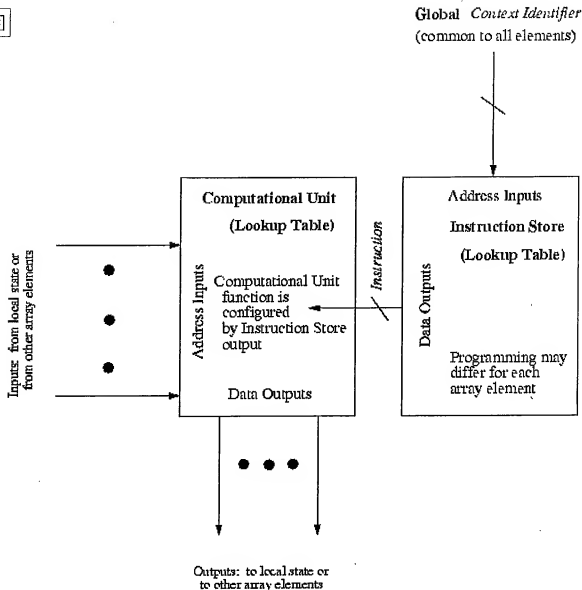




Lookup Table Model for DPGA



Figure  shows one way to view the DPGA AE. We think of the broadcast instruction as a *context identifier* (CID). The local *instruction store* in each computational element selects the executed instruction by table lookup using the CID as an address. Each AE holds a distinct instruction store, and hence different AEs execute distinct operations at the same time.





Configurable Instruction-Store View of DPGA AE

We can also think of the DPGA as an FPGA with multiple contexts. All the instructions programmed into a particular CID location in the AE's instruction stores can be thought of as one array-wide configuration. When we change the CID, we change contexts and instantly reconfigure the entire array.

In Figures  and , we only show how the CID selects the operation performed by the computational unit of each AE. In practice, the routing resources interconnecting AEs and the addresses of local stored data are also configured by the instructions. The CID must select the instructions which configure the interconnection and addressing as well as the computation.

The DPGA retains all the facilities of both FPGA and SIMD arrays. If we wish to perform traditional SIMD operations, we simply program the same instructions into all the instruction stores. If we wish to

perform traditional FPGA operations, we continually broadcast the same CID.

For maximum flexibility, the DPGA instruction data is writable by some combination of the instruction stream and local AE state. This allows the array to rapidly update context programming as necessary. If we think of the instruction store as a context cache, the writable instruction store allows high speed cache replacement. In cases where the sequence of distinct instructions is larger than the instruction store, we include code in the instruction stream to reload the instruction store.

Benefits

This hybrid architecture increases flexibility for both implementation and application of computational arrays. We can immediately identify several benefits which the DPGA architecture has over either SIMD arrays or FPGAs.

Implementation



The width of the CID field is unrelated to length of the stored instruction, but is typically much shorter. Since instruction distribution bandwidth is a major design constraint in high performance SIMD arrays, this reduction can be very important. For example, a complex FPGA cell might provide five input logic functions. Distributing a truth table for such a function requires 32 instruction bits; a CID as small as 8 bits might suffice for most applications, reducing the required instruction bandwidth by a factor of four.

A single DPGA array type can be engineered and produced to satisfy applications which have traditionally been served by either FPGAs or SIMD arrays. By combining the market base of these two array types, DPGAs can become commodity products, allowing cost effective implementation of novel data-parallel processors as well as more effective replacements for FPGA structures.

Applications

Multiple-Stream SIMD Array Processing

Viewing the DPGA from a traditional SIMD view, we have added the ability to simultaneously perform different operations. This is useful in cases where the pieces of the data in the array must be operated on differently or there is insufficient data parallelism to occupy the entire array performing the same computation on an operation cycle.

Boundary conditions are a common case where non-uniform data handling is required. In the DPGA model, boundary cells are programmed to perform a different set of operations from interior array cells during some computation. One particularly compelling example of such boundary processing is the use of a DPGA array for bit-parallel arithmetic operations, where the most- and least- significant bits of each parallel data word usually require special handling.

Configuration Cache and Time-Slice Computation



Viewing the DPGA as a traditional FPGA, we have added a configuration cache and the ability to perform zero-latency reconfiguration between cached configurations. With the hybrid configuration, we can perform computations in time-slices and achieve much more efficient utilization of AEs. In conventional FPGAs, most AEs are exercised at their full operational speed for only a small portion of

the clock cycle. That is, each AE is part of some path between sequential elements which are clocked at the component's system clock rate. Assuming there are \boxed{x} AEs in the critical path between sequential elements for a computation and each AE can operate in time \boxed{x} , the system clock cycles is at least \boxed{x} . A given AE only performs its function during \boxed{x} of the cycle and spends the rest of the time holding its value. By pipelining and time-slicing the computation, we can arrange for each AE to perform its computation at the correct time to produce the desired result without increasing the overall latency of computation. Each of the AEs is now available on the remaining \boxed{x} time slices to perform other operations. Consequently, we can schedule other operations on each AE during these unused cycles. We can effectively interleave computations and thereby extract higher throughput from an array of a given size.

In practice, each pipeline register introduces some timing overhead. If we call \boxed{x} the overhead of the pipeline register due to setup and hold times on the register, then the entire computation would really require \boxed{x} . If \boxed{x} , then the overhead is small and often worth paying to increase throughput. If \boxed{x} is larger, it often still makes sense to pipeline but at a larger granularity. In general, we can divide the \boxed{x} element critical path into \boxed{x} -element segments and only place pipeline boundaries between the segments. Here we slow the computation down as shown in Equation \boxed{x} .

$$\boxed{x} \quad \boxed{x}$$

At the same time, we allow an additional throughput as shown in Equation \boxed{x} .

$$\boxed{x} \quad \boxed{x}$$

Of course, traditional FPGAs can use their own registers to pipeline computations and increase throughput. The multiple loaded contexts allow each DPGA AE to be reused to perform different functions during each time slice.

Virtual Cells and Embedded Systems

In analogy to the Virtual Processor concept used in the SIMD Connection Machine [Hil85] to map a large computation onto a small number of processors, we can also treat the DPGA as having many *Virtual Cells* per physical AE. At a given point in time one function is active at each AE. Using the CID, we can switch the personality of each AE amongst the virtual cells it is emulating. This allows a small DPGA array to efficiently emulate a large FPGA array.

The virtual cells approach can also be viewed as a technique for reducing the system part count. Rather than collect enough FPGA components to spatially implement all the functions required at any point in time, we employ a single DPGA with an external memory or ROM chip to store additional contexts. A controller inside the DPGA sequences through contexts, swapping them from external memory when necessary. The DPGA can then switch between configurations sequentially to perform the complete

ensemble of required functions.

The logical extension of this integration process is to implement the DPGA in a high-density DRAM process and therefore integrate the programmable logic with a memory store. The combined chip is able to perform almost any logic operation, with a slowdown increasing with operation complexity.

Efficient Logic Emulation


FPGAs are in common use today for logic emulation. Most emulation systems directly map gates in the emulated system to gates in the FPGA. This uses the FPGA AEs inefficiently for the reasons identified above. The hybrid AEs can be used more efficiently by scheduling multiple gates in the emulated system to a single gate on the FPGA and using time-slice computation as above. The CID also provides the opportunity to perform *Event-Driven Simulation*. Like software oriented event-driven simulation, we can take advantage of the fact that some subsets of the system change infrequently. Using the virtual cells model, we can "swap" in a region of logic and emulate the region only when its value changes.

Processor Assistance

FPGAs also promise to be useful as tightly coupled co-processors for conventional microprocessors. The FPGA can be configured to perform some application-specific calculations more quickly than the processor [Sil93]. With the internal configuration cache, the DPGA array can switch between operations rapidly. This has two benefits. First, the array can support multiple configurations for a single computation thread. Second, configuration contexts allow DPGA co-processors to support multiple threads. This capability will become more important for fine-grained, multithreaded microprocessors which support fast context switching (e.g. April [ALKK90], *T [NPA92]).

Costs

From both the FPGA and SIMD array standpoint, the primary additional cost for the DPGA array is the

area for the instruction store lookup table (See Figure ). The instruction store can be implemented with a single-port, SRAM-style memory array which can be implemented very compactly in a full-custom design. For example, 70% of each Abacus AE [Bo93] comprises roughly 80 bits of SRAM memory. The Abacus AE requires 16 instruction bits to specify the operation performed by the ALU on each cycle. We could add another 80-bit SRAM to serve as a 5-entry context cache. If we distribute a fully decoded CID to avoid the area overhead of a decoder at each instruction store, we reduce the instruction distribution bandwidth from 16 bits per cycle to 5 bits per cycle while increasing the AE cell area by 70%. The area cost may not even be an issue if the number of AEs on a chip is limited by the number of pins available.

Indirection through the instruction store lookup array need not impact the AE cycle time. We can add a pipeline register between the program store output and the array computational element. This pipelining allows the array computational element to run just as fast as it did in the SIMD or FPGA array. The addition of a pipeline stage adds another latency stage to instruction distribution. This has no net loss when compared to the FPGA case where no instruction distribution occurs. The additional cycle of latency due to this pipeline stage is small compared to the typical latency in the instruction distribution path for a SIMD array. In a modern high-speed (> 100 MHz) SIMD array, there are at least three pipeline stages in the instruction delivery phase. Further, there are at least three stages in the return conditional information path. For efficient operation, SIMD arrays generally run long instruction sequences without interruption so that pipeline depth is not a significant issue.

Previous SIMD Approaches to Local Configuration

The need for local configuration in SIMD arrays is demonstrated in part by some small steps which SIMD architectures have taken towards local configuration. Several architectures allow local state bits to modify the transmitted instruction, the transmitted address, or the local network connections. This local modification is referred to as *operational*, *address*, and *connection autonomy* where autonomy indicates that processors have the ability to not perform identical operations [ML89]. The translation process is typically very simple: operational autonomy is achieved by potentially inverting the ALU output; addressing autonomy by modifying a few low order bits of the address [BDHR88]; connection autonomy by configuring a network crossbar at each AE site [ML89].

All of these techniques can be viewed as resource-limited implementations of the DPGA architecture. Each can be derived by starting with a small number of contexts and observing that the instructions in each context are almost identical except for a few bits. Factoring out the common bits in hardware leads immediately to the *ad hoc* implementations described above. Thus, the DPGA model subsumes the local configuration mechanisms in SIMD machines.

Array Specialization and Optimizations

In this section we show that the different computing styles used by the FPGA and SIMD architectures drive important architectural trade-offs. We also show that these trade-offs have led to architectural extremes which are often not optimal given available technologies. We show examples where improvements result when we match architecture to the available technology. In many cases, technology-oriented optimizations move us away from the architectural extremes, increasing the similarity between the resulting SIMD and FPGA arrays.

Architectural Implications of Computing Style

Most implementations of the two architectures compute in two distinct styles: FPGAs unroll their computations in space, while SIMD machines unroll their computation in time. This unrolling can be described as bit-parallel and bit-serial computations, respectively. Bit-serial techniques exhibit comparatively high latency, but are very efficient in terms of throughput and silicon area; bit-parallel techniques produce results quickly but require considerable silicon and routing resources. Because most FPGAs were used to construct controller circuits in which the operation latency was crucial, they were optimized for bit-parallel operation. SIMD AEs were optimized for high throughput, and therefore for bit-serial operation. Upon examination of the two computing styles, the architectural differences required to support them become obvious. The three key areas for comparison are interconnect resources, local state, and clocking strategy.

Bit-parallel operation requires significant interconnection resources to compose complex logical functions of simultaneously existing bits, and the interconnect pattern need not change through time. Local state is largely unnecessary since intermediate result bits are stored on the wires. Comparatively slow clocking is required to allow computation to ripple through several combinational units in a single cycle.

Bit-serial operation has the opposite requirements. Only a simple logical function, usually on the order of a one-bit adder, need be performed at each step. Some local state is required to store the intermediate results. Little wiring is required since the composition occurs in time, not in space. In some sense,

memory cells replace the wires of a bit-parallel design for both storage and communication. Finally, fast clocking is desirable since many simple steps are required for any significant computation.

It is instructive to note that when an implementation of each architecture adopts the opposite computing style, it begins to resemble the opposite architecture. For example, the Abacus machine, a SIMD architecture designed for bit-parallel operation, has a four-input/two-output logical unit, two wires, and only 64 state bits in each AE [Bo93]. In the opposite direction, the Concurrent Logic FPGA AE, targeted specifically for computing rather than implementing random logic, has reduced routing resources, a simple logical unit, and is intended to be clocked at high rates [Fur93].

In the remainder of this section, we discuss these three differences in further detail. We note how the available technology influences implementations. In many cases, technology encourages implementations of these arrays which are not as disparate as the architectural ideals.

Interconnection Style

The main difference in interconnection between the SIMD and FPGA architectures is that the first dynamically routes multiple logical signals over a single wire, while the second statically assigns a dedicated wire to each signal. In the following two examples, the adoption of the opposite communication style has improved performance.

Virtual Wires

Since FPGAs are notoriously short of input/output pins, a common problem with FPGA implementations is the difficulty of partitioning a design among multiple FPGA components. As a result of a limited number of I/O pins, only a small fraction of the AEs in each FPGA component is generally usable. A recent technique called *virtual wires* [BTA93], has been used to overcome this limitation by sending different signals through the same pin at different points in time. This is another name for time-multiplexing I/O pins, an approach employed by SIMD computers for many years. The improvement due to this technique is possible because the I/O bandwidth was underutilized when each I/O pin was statically assigned to a single logical signal. Time-multiplexing allows the array to use the additional available bandwidth to better balance internal silicon utilization with cross-chip signalling bandwidth.

Static routing, as practiced in the SIMD community, may well be applicable to better utilization of internal routing resources. Chip area allocated to programmable interconnect may be reduced if the wires can be used more efficiently by high-speed, time-multiplexed operation.

Multiple Networks

All SIMD array architectures to date have used a grid with a single wire between AEs, while FPGA AEs typically have access to at least four wires. The performance of SIMD machines, especially of the software bit-parallel [Bo93] variety, can be increased by adding extra intra-chip grid wires.

Multiple wires are possible in many SIMD implementations because local routing channels to route additional signals between AEs take up little or no additional area in the SIMD array. The additional channels can be used to increase the communication bandwidth between AEs and therefore increase computational efficiency.

Specifically, multiple wires reduce the overhead of dynamically reconfiguring the SIMD network. For example, in a bit-parallel computation the result of a data word comparison must be broadcast to each

AE operating on a bit of the data word, so that the entire group can be disabled or enabled as a unit. But in a magnitude comparison, the result bits travel from LSB to MSB. Upon completion, the bit stored in the MSB must be broadcast, which involves turning the network around. This turning around step can take almost as long as the actual comparison. With two network wires, the overhead is eliminated completely.

Clocking Strategy

Modern SIMD architectures are designed with a very high clock rate to compensate for the simple logic operation in each cycle. In contrast, FPGA arrays are clocked comparatively slowly, to allow propagation through several stages of combinational elements. We now describe why FPGA arrays should be clocked much *faster* than SIMD arrays.

In the current technology trend, on-chip clock speed is increasing much faster than off-chip communication bandwidth. This presents a problem for SIMD arrays. For example, the internal cycle time of a SIMD AE implemented in a 0.8 micron CMOS process can be as low as 5 nanoseconds. Delivery of a wide, 50-bit instruction at 200 MHz to each SIMD chip is a challenging design problem. We see that the factor limiting cycle time is the instruction bandwidth and that ultimately, SIMD AE complexity will have to increase in order to match internal cycle time to available bandwidth.

FPGA architectures do not suffer from these bandwidth limitations since they do not require instructions, and can therefore only benefit from switching to high-speed clocked, pipelined operation. External system design complexity need not increase, as the fast clock can remain purely internal, with chip interfaces operating at low speed. The fast internal clock can be generated with a phase-locked loop from a slow external clock [HCC+93]. Also, a clocked system allows the use of low-power, clocked logic families such as dynamic logic.

Local State

As discussed earlier, SIMD machines require considerable local state to store intermediate results, while FPGAs store these results on the wires connecting combinational blocks. The problem with a large amount of state is the need to address it. For example, most SIMD AEs have at least 64 local state bits, requiring 6 bits to specify each read and write address. This is significantly more bits than the 8 required to fully specify a 3-input ALU operation. Clearly, the need to address local state can dominate the instruction bandwidth requirements of SIMD architectures.

FPGA AEs do not have this problem, as their local state typically consists of a single register, which is addressed implicitly by connecting it to the logical unit. In practice, some FPGA cells can be used as a small memory, but we do not consider this part of an AE, as many AEs must be connected together to implement an address generator.

This dissimilarity is another aspect of the static versus dynamic allocation encountered in the interconnect control. When FPGAs need to access a state bit, they statically allocate a wire from the memory cell to the appropriate logic block; SIMD machines must pay the overhead of dynamic access on every cycle, even if the same address is being used for many cycles in a row. Once this observation is made, several techniques for reducing the overhead of dynamic addressing can be evaluated with respect to a particular technology. These techniques include implicitly addressed accumulators, register windows, and register renaming.

<input type="checkbox"/> Characteristic	FPGA	SIMD
Wires/direction	4-8	1
Routing	Static	Dynamic
Logical unit	Logical function of 4-8 inputs	1-bit adder
Clock	Slow	Fast
Arithmetic	Bit-parallel	Bit-serial
Local state	1-2 bits	64-256 bits

Conclusions ☐

Starting from a common computational array model, we examined the similarities and differences between FPGA and SIMD arrays. While both are composed from arrays of fine-grained computational elements, they differ substantially in how they compute. SIMD arrays are employed for high-throughput computations using bit-serial computing techniques. To match this function, SIMD arrays are optimized to vary computation in time while performing the same operation on all array elements at each time step. FPGAs are employed for low-latency computation using bit parallel computing techniques. Optimized for bit-parallel computations, FPGAs are configured statically, and perform computation spatially.

We introduced a hybrid computational array architecture, the DPGA, and showed that it can provide better performance than either. The DPGA allows computation to vary both spatially and temporally. Using a local instruction store, the DPGA performs spatially varying computation requiring no additional bandwidth. The DPGA mixes both bit-parallel and bit-serial computations in a single array structure. We suggested examples where this additional flexibility allows higher performance or lower part count than pure FPGA or SIMD alternatives.

Finally, we explored the influence of bit-serial and bit-parallel computing styles on the details of array implementation. We saw that these two computing paradigms push array organization towards opposite extremes of clocking, routing, and state management. We further saw that these extremes are often less than optimal when one considers the technology available for array implementation. Hybrid approaches often extract the highest performance in a given technology. Based on these observations, we believe there is room for significant cross-fertilization of ideas between the FPGA and SIMD communities.

See Also...

- DPGA-coupled Processor
 - [\[tn100 HTML link\]](#) [\[tn100 FTP link\]](#)

References

Act90 Actel Corporation, 955 East Arques Avenue, Sunnyvale, CA 94086. *ACT Family Field Programmable Gate Array DATABOOK*, April 1990.

ALKK90

Anant Agarwal, Beng-Hong Lim, David Kranz, and John Kubiawicz. APRIL: A Processor Architecture for Multiprocessing. In *Proceedings of the 17th International Symposium on*

Computer Architecture, pages 104-114. IEEE, May 1990.

BBLC93

M. Bolotski, R. Barman, J. J. Little, and D. Camporese. SILT: A Distributed Bit-Parallel Architecture for Early Vision. *International Journal of Computer Vision*, IJCV-11:63-74, 1993.

BDHR88

D. W. Blevins, E. W. Davis, R. A. Heaton, and J. H. Reif. Blitzen: A Highly Integrated Massively Parallel Machine. In *Frontiers of Parallel Computation '88*, 1988.

Bla90 Tom Blank. The MasPar MP-1 Architecture. In *International Conference on Computer Architecture*, 1990.

Bol93 Michael Bolotski. Abacus: A High-Performance Software Bit-Parallel Architecture. AI memo, in preparation, MIT, 1993.

BRV93

Patrice Bertin, Didier Roncin, and Jean Vuillemin. Programmable Active Memories: A Performance Assessment. In *Research on Integrated Systems Symposium*, 1993.

BTA93

Jonathan Babb, Russell Tessier, and Anant Agarwal. Virtual Wires: Overcoming Pin Limitations in FPGA-based Logic Emulators. In *FPGA '93*, 1993.

Fur93

F. Furtak. A Field-Programmable Gate Array for Systolic Computing. In *Research on Integrated Systems Symposium*, 1993.

Gea91

Maya Gokhale and et al. Building and Using a Highly Parallel Programmable Logic Array. *IEEE Computer*, 24(1):81-89, January 1991.

HCC +93

Mark Horowitz, Andy Chan, Joe Cobrunson, Jim Gasbarro, Thomas Lee, Wing Leung, Wayne Richardson, Tim Thrush, and Yasuhiro Fujii. PLL Design for a 500MB/s Interface. In *ISSCC Digest of Technical Papers*, pages 160-161. IEEE, February 1993.

Hil85 D. Hillis. *The Connection Machine*. Distinguished Dissertations. The MIT Press, Cambridge, Massachusetts, 1985.

Inc91 Xilinx Inc. *The Programmable Gate Array Data Book*. Xilinx, Inc., 1991.

Lea91

R. M. Lea. Wafer-Scale Massively Parallel Computing Modules For Fault-Tolerant Signal and Data Processing. In *International Conference on Computer Architecture*, 1991.

ML89

M. Maresca and H. Li. Toward connection autonomy of fine-grain SIMD parallel architecture. In *Parallel Processing for Computer Vision and Display*, pages 77-86. Addison-Wesley, 1989.

NPA92

Rishiyur S. Nikhil, Gregory M. Papadopoulos, and Arvind. *T: A Multithreaded Massively Parallel Architecture. In *Proceedings of the 19th International Symposium on Computer Architecture*. ACM, May 1992.

Pot85 Jerry L. Potter, editor. *The Massively Parallel Processor*. MIT Press Series in Scientific Computation. The MIT Press, Cambridge, Massachusetts, 1985.

Qui92

QuickLogic Corporation, 2933 Bunker Hill Lane, Santa Clara, CA 95054. *QuickLogic: Very High Speed FPGAs*, 1992. QuickLogic is now a subsidiary of Lattice Semiconductor Corporation.

Sil93 Harvey F. Silverman. Processor Reconfiguration Through Instruction-Set Metamorphosis. *IEEE Computer*, 26(3):11-18, March 1993.

MIT Transit Project